

A Neuro-evolutionary Approach to Control Surface Segmentation for Micro Aerial Vehicles*

Max Salichon
Oregon State University
max.salichon@gmail.com

Kagan Tumer
Oregon State University
kagan.tumer@oregonstate.edu

Abstract

This paper addresses control surface segmentation in Micro Aerial Vehicles (MAVs) by leveraging neuro-evolutionary techniques that allow the control of a higher number of control surfaces. Applying classical control methods to MAVs is a difficult process due to the complexity of the control laws with fast and highly non-linear dynamics. These methods are mostly based on models that are difficult to obtain for dynamic and stochastic environments. Moreover, these problems are exacerbated when both the number of control surfaces increases and the model's accuracy in determining the impact of each control surface decreases. Instead, we focus on neuro-evolutionary techniques that have been successfully applied in many domains with limited models and highly non-linear dynamics. Wind tunnel simulations with AVL show that MAV performances are improved both in terms of reduced deflection angles and reduced drag (up to 5%) over a simplified model in two sets of experiments with different objective functions. We also show robustness to actuator failure with desired roll moment values still attained with failed actuators in the system through the neuro-controller.

Keywords: Evolutionary algorithms; Micro Aerial Vehicles, Neural Networks.

1 Introduction

Micro Aerial Vehicles (MAVs) provide an appealing solution to many tasks such as monitoring [19, 18], reconnaissance, sensing [12], search and rescue, and enemy targeting [24] that may not be accomplished by larger Unmanned Aerial Vehicles (UAVs). MAVs can accomplish such demanding missions without posing a large threat in case of malfunction, since they are by definition small and light. However, MAVs are notoriously difficult to control, and recent work is focusing on developing stable flight characteristics [13, 32, 33, 15, 22, 10].

In general, MAVs have limited power, limited computational power and limited control surfaces and actuators, which makes their control and navigation a considerable technical challenge. In addition, because of their size and weight, they operate in regimes that are highly non-stationary and have highly non-linear dynamics, making their control even more difficult. MAVs range in size from 6 to 24 inches (15 to 60 cm), and fly at speeds of 10 to 50 mph (5 to 20 m/s) [1, 14]. In addition, MAVs must be able to operate around buildings and

*Appears in the International Journal of General Systems, Vol 42, No. 7, pp. 793-805, 2013.
DOI:10.1080/03081079.2013.776203

obstacles both in and outside of cities. As such, they need to be highly maneuverable and have some robustness to changing environmental conditions (e.g., wind gusts) [5, 26].

Recent work has also focused on flexible-wing MAV designs to improve both the stability of the vehicle and its robustness to wind gusts. These performance gains are the result of the wing deforming continuously and absorbing the energy created by the instabilities of the air flow rather than transmitting all that force to the motion of the MAV [1]. Indeed, flexible wings can lead to a higher airspeed, higher climb rate, improved maneuverability, and a higher lift to drag ratio which is particularly important for MAVs as it improves their gliding capabilities.

The key performance improvements for flexible-wing MAVs stem from the gains in energy absorption characteristics of the deformable wings. In addition to this passive control benefits, one can use an actuator to change the shape of the wing during flight where roll control of an MAV was achieved by actively morphing the wing [10]. A first step towards an actively deformable wing is one with segmented control surfaces which is a more practical solution both in terms of simulation and fabrication. This approach was implemented on a remote controlled 5.5ft wingspan UAV [2]. The wing ailerons were divided into 16 independent control surfaces that each had their own actuator. A reconfigurable controller was developed to actuate all 16 servos depending on the configuration. Flight tests showed promising results and improved performance over the unmodified aircraft. Those tests demonstrated the concept of segmented control surfaces and provided good preliminary results but provided no method for finding an optimal actuation mode for the system.

A multiagent controller for segmented control surfaces was also proposed [6, 7] which aimed to determine the impact of each control surface on the UAV performance using an objective decomposition approach [30, 34]. The resulting system used the concept dubbed “MiTE” which stands for Miniature Trailing Edge Effector, which are devices actuated with a deflection angle of up to 90 degrees and are 1-5% of the chord in height. The UAV used for the experiment was a flying wing with 6ft wingspan and 30 degrees of leading edge sweep. The results showed the promise of using multiagent control for UAVs [6, 7].

Evolutionary computation techniques have been used successfully to solve benchmark control problems including the inverted pendulum [21] and the ball and beam [16] problems. In addition, they have been used in real world applications such as searching through a space of plans generated from a planning algorithm to yield good control policies in a planetary rover control problem [9] and using “sub-populations” to control rockets [11].

Furthermore, evolutionary algorithms have also been extended to complex control problems such as multi-rover control problems [3, 31] where a large number of agents have to maximize the overall system objective as well as their own objective [20, 27, 23]. This technique is therefore well suited for controlling an MAV with segmented control surfaces. For example, having agents try to maximize a global evaluation function through a process of finding good collaborators that avoid suboptimal equilibria [20] or scaling evaluation functions to ensure that the agents do not work at cross-purposes have been explored [25].

Contribution of this Work

In this paper, we show that (i) neuro-evolutionary techniques can be used to control multiple surfaces to improve the flight characteristics of an MAV by designing appropriate objective functions (e.g. roll moment value); and (ii) the neuro-controller is robust and can achieve the desired roll moments even in the presence of failures. Section 2 describes the platform and experimental setup, Section 3 shows the experimental results where drag on the MAV was reduced by up to 5%, Section 3.4 shows the results with failures present in the system, and

Section 4 discusses the relevance of the results and highlights directions for future work.

2 Segmented Surface Control for MAVs

The key in determining how to control the segmented surfaces of an MAV is in both modeling the vehicle well and establishing the correct objective functions. In this section, we first present the specific MAV configuration, then provide three different objective functions for assessing performance, and finally discuss the approaches to optimizing those functions.

2.1 Micro Aerial Vehicle Characteristics

In this work, we use a modified GENMAV [26], an MAV developed by the Air Force Research Laboratory Munition Directorate (AFRL/RW). This MAV include a 24 inch wingspan with a 5 inch chord, circular fuselage 17 inches long, and a dihedral angle of 7 degrees. This MAV was designed for a flight speed of between 10 and 50 mph with an average flight speed around 30mph. We use the vortex-lattice method aeroprediction code AVL (Athena Vortex Lattice) to compute the aerodynamic characteristics of this MAV [26].

In these experiments, we used a modified GENMAV with a different wing for improved low speed performance as well as a conventional tail (as opposed to the V-tail) to include a greater number of control surfaces. In addition, we further modified GENMAV to include a greater number of control surfaces. As a first step, only the tail sec As a first step, only the tail section was modified with the elevons broken down into multiple control sections. Test configurations include 4 elevons on each side of the tail, for a total of 8 elevons. A configuration with 12 elevons was also tested (6 elevons on each side of the tail).

2.2 System Objective Functions

The selection of objective functions play a key part in determining the success of any control algorithm. The key in for the MAV task is to both meet the target value of the desired forces/moments and minimize the actuator angles of the different control surfaces of the elevator (8 and 12 controls surfaces total).

2.2.1 Minimizing the Drag

The first objective function G_{F_d} (See Results in Section 3.1) is calculated using the drag and roll moment desired value.

$$G_{F_d} = \alpha C_1 \left[C - \left(\frac{L_a - L_d}{L_d} \right)^2 \right] + (1 - \alpha) C_2 [F_{d_M} - F_d]^2 \quad (1)$$

Where C , C_1 , and C_2 are normalization constants with values of 9, 200, and 484000 respectively, F_d is the force of drag, F_{d_M} is constant and equal to 0.142, L_d and L_a are the desired and actual roll moment values. In this case, a good value for α is 0.998. The drag calculation is done internally in AVL.

2.2.2 Minimizing Elevon Angles

The second objective function G_{ω_1} (See results in Section 3.2) used is calculated using the actuator angles and roll moment target value.

$$G_{\omega_1} = \alpha C_3 \left[C - \left(\frac{L_a - L_d}{L_d} \right)^2 \right] + (1 - \alpha) C_4 \sum_{i=1}^N [\omega_M - |\omega_i|]^2 \quad (2)$$

Where L_d and L_a are the desired and actual roll moment values, ω_i is the deflection of the control surface (with a maximum deflection of $\omega_M = \pm 30$ degrees for each actuator), and C , C_3 and C_4 are normalization constants with values of 9, 200, and $2/N$ respectively, and N is the number of control surfaces. For these experiments, α needs to be 0.998 or above, otherwise the roll moment target value cannot be reached.

2.2.3 Minimizing Relative Angle Between Elevons

Another similar objective function G_{ω_2} that was used was designed to minimize the actuator deflection angles relative to each other.

$$G_{\omega_2} = \alpha C_5 \left[C - \left(\frac{L_a - L_d}{L_d} \right)^2 \right] + (1 - \alpha) C_6 \sum_{i=1}^{N-1} \left[\omega_M - \frac{|\omega_i - \omega_{i+1}|}{2} \right]^2 \quad (3)$$

Where C , C_5 , and C_6 are normalization constants with values of 9, 200, and $2/N$ respectively, L_d and L_a are the desired and actual roll moment values, ω_i is the deflection of the control surface (with a maximum deflection of $\omega_M = \pm 30$ degrees for each actuator), and N is the number of control surfaces.

2.3 Optimizing MAV Objectives

In this work, we use a feed-forward neural network trained using a neuro-evolutionary algorithm ([3, 11, 17, 28]) to optimize the MAV objective functions discussed above. The neural network learns the configuration of the control surfaces through the system objective function that is designed to achieve the targeted forces and moments and at the same time minimize the control surfaces' deflections or drag.

This algorithm employs an evolutionary process that starts with a set of randomly initialized neural networks. New neural networks are generated by modifying (mutating) parameters using values sampled from a Cauchy distribution. In subsequent time steps (generations) the algorithm selects a candidate controllers from the pool using ϵ -greedy selection. The new network is stored in the pool only after an agent has used it and sampled their resulting performance, with the poorest performing network begin discarded. The population size was set at ten networks.

The structure of the neural networks consisted of a feed forward neural network with one hidden layer [8]. This network has six inputs (corresponding to the total forces and moments applied to GENMAV), and eight or twelve outputs (corresponding to the angles of the elevator control surfaces for). In these experiments with eight and twelve control surfaces, each control surface could move independently between -30 and +30 degrees. The results reported below are based on neural networks of twelve hidden units (the results are fairly insensitive to this parameter), an epsilon-greedy selection probability of $\epsilon = 0.05$, a level of initial weights of $\gamma = 0.1$, a level of mutations of *mutate* $\gamma = 0.05$, and a probability that a weight will be mutated of 0.02. These parameters were then kept constant for the experiments described in section 4. All results are based on $N=12$ runs and for a standard deviation of σ , we show the differences in the mean ($\frac{\sigma}{\sqrt{N}}$) when appropriate.

Finally, we coupled the neural networks to the Athena Vortex Lattice (AVL) software package which is an aerodynamic prediction code based on a vortex-lattice method. AVL was used to estimate the aerodynamic characteristics of GENMAV under different conditions and configurations. The output of AVL includes the forces and moments for the entire configuration as well as the lift and drag coefficients.

The simulation runs consist of providing forces and moments as inputs to the neural network, obtaining elevon angles from its outputs, running AVL to provide the resulting aerodynamic parameter values, computing the objective function, and having the neural network learn from the objective function.




Figure 1: GENMAV in AVL

3 Experimental Results

In order to evaluate the impact of using multiple control surfaces and training a neuro-controller to optimize the control surface angles, we performed the following experiments:

- The basic configuration consisted of a GENMAV with 2 elevons. This was used to obtain the aerodynamic parameter values that are used as a reference.
- A neuro-controller was used to control segmented control surfaces to explicitly minimize the:
 - drag (Section 3.1)
 - actuator angles (Section 3.2)
 - relative actuator angles (Section 3.3)
- Finally, a neuro-controller was used to control the system in the event of an actuator failure.

3.1 Neuro-Controller Evolved to Explicitly Minimize Drag: G_{DRAG}

We explored the potential to explicitly reduce drag by incorporating a drag term in the objective function that the neuro-controllers aims to optimize. This objective function (described in Equation 1) directly accounts for drag and roll moment target values but does not include the actuator angles.

Figures 2 and 3 presents the results for the drag data with G_{DRAG} . Figure 2 shows an example of the elevon positions for the configuration with 2 and 8 elevons. While these solutions are similar to the ones found with G_{DEF} and G_{DEF2} , the solution provided by G_{DRAG} shows more symmetry in the elevon configuration.

Figure 3(b) shows the results for the drag with the 3 different objective functions G_{DRAG} , G_{DEF} , and G_{DEF2} . G_{DEF2} produces results that are similar to G_{DRAG} , which indicates that minimizing the relative angles in between elevons can be used to indirectly minimize the drag. This is a particularly important result since the drag was available through the use of AVL, the aero-prediction code each time the elevon configuration is modified. The drag calculation for those configurations takes a significant amount of time and is usually not available directly when using flight simulators, and would not be possible for real flights of an MAV platform. Minimizing the relative deflections of the elevons can therefore provide a very good alternative to using the drag directly in the objective function calculations.

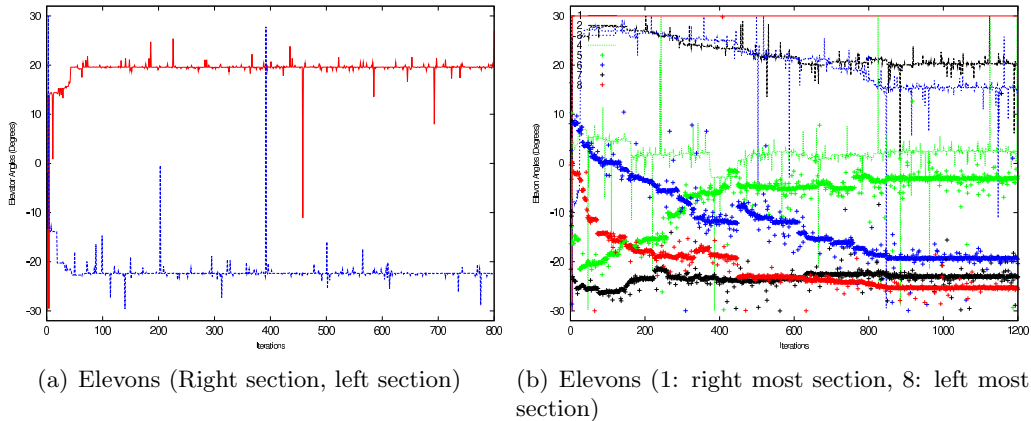



Figure 2: Elevon Angles with 2 and 8 control surfaces (Min Drag, Roll Moment = 0.028)

This solution is the intuitive solution that we would expect when trying to induce roll on the MAV while trying to minimize drag.

3.2 Neuro-Controller Evolved to Minimize Actuator Angles: G_{DEF}

Minimizing the elevons angles provides improved MAV flight characteristics such as smoother flight maneuvers which is an important benefit for MAVs. This section shows the results of experiments where several roll moment target values are achieved while at the same time the actuator angles are minimized.

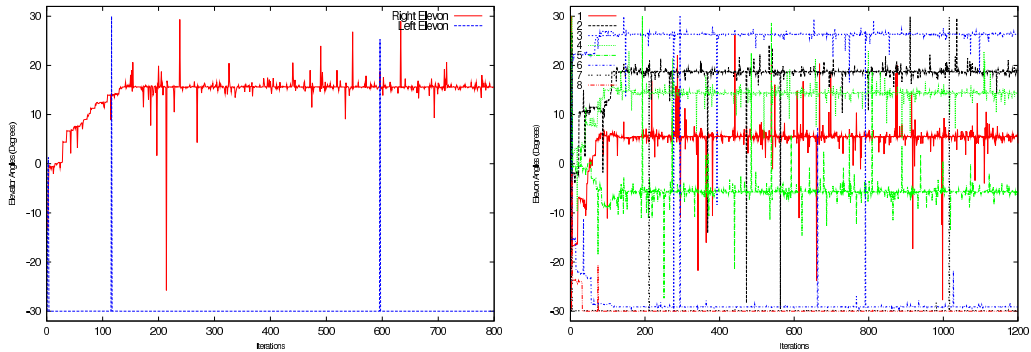
Figure 4 shows an example of the elevon angle values for a target roll moment value of 0.030. The elevon angles are progressively minimized as the neural network learns the optimal solution for the desired roll moment. To achieve a desired roll moment value of 0.030 with the standard configuration (2 elevons) requires the right and left elevons to move to 15 and -30 degrees respectively as shown in Figure 4(a). Figure 4(b) shows the elevon angles with the 8



(a) Drag vs Roll Moment (2, 8 and 12 Ctrl Surf) (b) Drag between the three objective functions (2 ctrl surf)

Figure 3: Drag Results

control surfaces MAV configuration. This configuration reduces the elevon angles which allows for smoother maneuvers and does not require as much effort from the actuators.



(a) Elevons (Right section, left section) (b) Elevons (1: right most section, 8: left most section)

Figure 4: Elevon Angles with 2 and 8 control surfaces (Roll Moment = 0.030)

A second and arguably more important benefit of segmented control surfaces is the potential for drag reduction. Figure 5(a) shows the drag results for the MAV with 8 elevons versus the MAV with its original configuration (right and left elevons). This particular MAV configuration coupled with the first objective function G_{DEF} (Section 2.2.2) does not exhibit any significant drag reduction suggesting that another objective function might provide better results. Another intuitive solution would be to minimize the relative angle between an elevon and its two direct neighbors. This solution is presented in Section 3.3.

3.3 Neuro-Controller Evolved to Minimize Relative Actuator Angles: G_{DEF2}

Figure 6 show similar results as Section 3.2 with the second objective function G_{DEF2} presented in Section 2.2.3. The elevon angles are minimized for smoother flight maneuvers. As with the first objective function G_{DEF} , no significant drag reduction can be observed between the configuration with 2 and 8 elevons (Figure 7(a)). However, G_{DEF2} induces significantly less drag than G_{DEF} for some of the lower values of the roll moment as can be seen in Figure

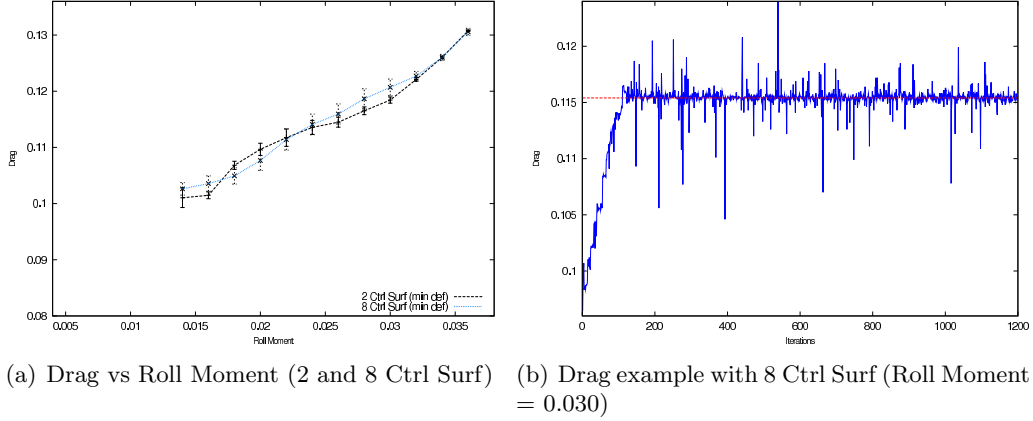


Figure 5: GENMAV Drag Data (Minimize Actuator Angles)

7(b). G_{DEF2} would therefore be a better objective function than G_{DEF} because it effectively minimizes the elevon angles while at the same time inducing less drag.

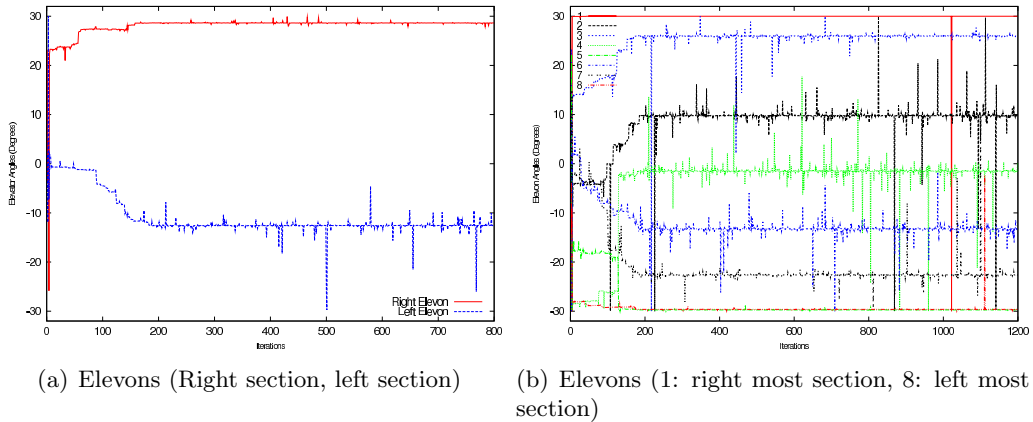


Figure 6: Elevon Angles with 2 and 8 control surfaces (Roll Moment = 0.028)

The results presented in this paper are promising and show the potential for improving MAV flight characteristics and increasing MAV robustness by using a larger number of control surfaces. The control of such a modified MAV is possible with the use of a neural network that if properly tuned and trained can provide optimal solutions to the MAV control problem.

3.4 Neuro-controller for Actuator Failure

The graphs in Section 3.4 shows that it is possible for a Neural Network to learn and adapt to changes in the environment, in this case failure of an actuator which changes the system's dynamics in order to regain control of the MAV. First Figure 8 shows that the desired roll moment can still be achieved when an actuator fails by finding a new solution that compensates for that failure. This is the case for both objective functions, minimize the elevon angles and minimize the relative deflections between elevons.

Figure 9 shows the drag results when comparing the system with and without failures. Results show that the drag is not negatively impacted when a failure occurs in the system.

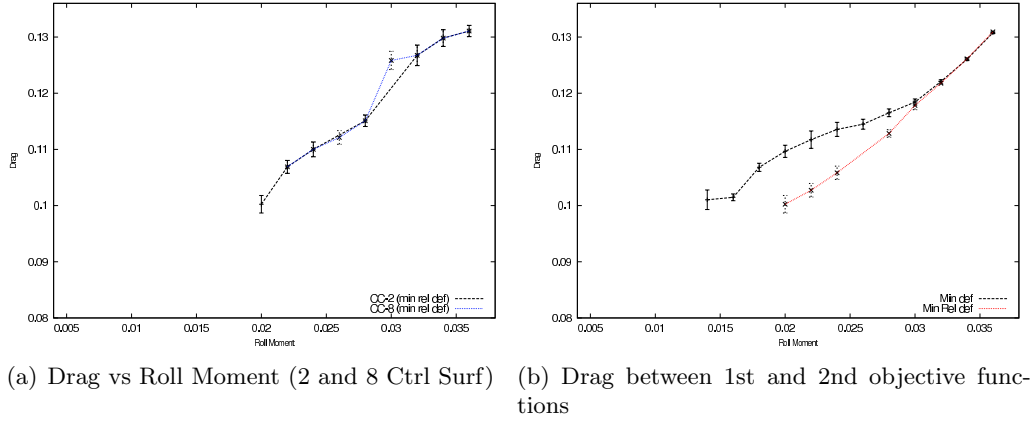


Figure 7: Drag Results (Minimize Relative Deflector Angles)

Drag is even slightly better when a failed actuator in some cases depending on how and where the failure occurs. This is primarily due to the objective function that in these cases minimize the actuator angles or relative angles between actuator which indirectly reduces the drag but does not necessarily always find the optimal solution for minimizing the drag.

These results demonstrate that control surfaces segmentation can be controlled by a neuron-evolutionary based controller in the event of an actuator failure therefore increasing the robustness of the platform. These results are important and show that a neuro-controller trained to the specifics of an MAV in simulated flight combined with control surface segmentation could be used for critical missions where actuator failures could be managed so that the mission can still be completed. To be applicable on an actual platform, it is important to note that actuator failures need to be detectable by the system through some type of sensing mechanism.

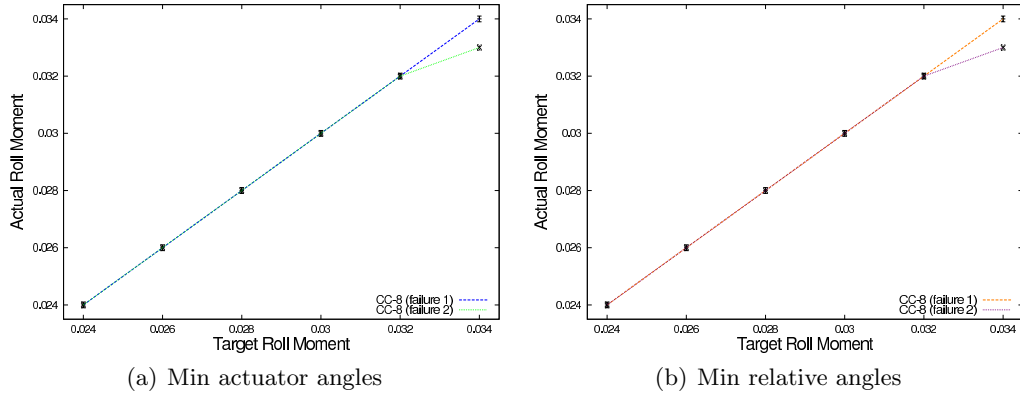


Figure 8: Target vs Actual Roll Moment when failures occur

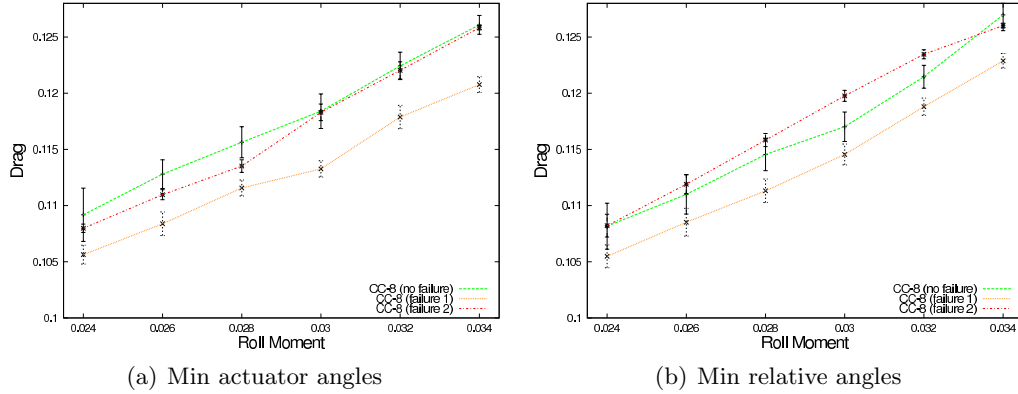


Figure 9: Drag Results: Failures

4 Discussion and Future Work

MAVs provide both the promise of new data gathering tools and the challenge of a difficult to control systems. To date, their applicability to the domains in which they are the most needed (e.g., dangerous search and rescue or reconnaissance) has been limited by the difficulties in obtaining good control algorithms. This paper presents control surface segmentation as a novel approach to the MAV control problem. In addition, it provides improvement of the flight characteristics by introducing a larger number of control surfaces on the elevon control sections.

Sections 3.2,3.3 and 3.1 showed the effectiveness of using neuro-evolutionary controller for segmented control surfaces. Using segmented control surfaces allows for smoother flight characteristics and flight maneuvers through minimization of actuator angles. Additionally, drag reduction of up to 5% can be seen for the larger values of the roll moment. If drag reduction is the objective, and if the drag is directly available, direct use of the drag in the objective function calculations provides the best results. However, if the drag is not available, minimizing the deflection between elevons still provides similar results and could be used instead.

Results showed a drag improvement of up to 5% which was obtained by a gradual deflection of each actuator which lead to a smoother control effort. Also, simulations conducted with 8 and 12 elevons showed no significant differences between the 2 configurations which indicates that for this particular problem and configuration, 8 elevons are sufficient and increasing the number of elevon segments will not improve the drag or efficiency of the MAV. The solutions provided by the neural network matches the intuition that a gradual actuator deflection would provide close to optimal solutions. Results presented in this paper show the potential of such configurations to improve flight characteristics of MAVs that are inherently difficult to control. Neural networks can effectively learn from the system and provide an optimal system's configuration therefore allowing such modifications on an MAV platform. Furthermore, such a configuration would provide a higher level of robustness to the system that could recover and adapt from potential failures of some elements in the system which is critical for completing the assigned missions as seen in section 3.4. The Neural Network was able to learn and adapt to the new MAV configuration that included a failed actuator and was able to provide a new solution for the control strategy in order to stay in control of the vehicle.

The results presented in this paper are a first step that shows the potential of leveraging

learning methods to accommodate a larger number of control surfaces on an MAV. Using these methods allows improvements in the flight characteristics of MAVs as well as provide more robust control strategies where recovering from potential failures is critical. More experiments will be conducted using a similar configuration to improve upon the results shown in this paper. Such experiments could include fine tuning the system objective function to improve the drag reduction in different situations and for different desired values of some parameters. Another valuable experiment would be to repeat similar experiments with two or more different algorithms such as Q-Learning to see if one would perform better than the other for different configurations. A significant improvement is expected with the use of multiagent techniques applied to the MAV control problem [3, 29, 30]. The system would then consist of independent agents (control surface actuators) that would learn to maximize a reward that would be specific to each agent but that would benefit the overall system.

Another important goal is to provide a flight controller that increases the robustness of the MAV to wind gust and various perturbations in simulated and real flight conditions. This is achieved with a flight dynamic simulator (JSBSim) that provides changes in the environment so that different control strategies can be established to maintain control of the vehicle in situations where PID based controllers do not perform well due to the high instabilities. The control loop consists of the flight simulator advancing the state of the system to the next time step using the neuro-controller's outputs (elevator position, motor speed), the neuro-controllers then uses the new states information as inputs in order to provide the next control inputs to the flight simulator. Robustness to wind gusts and perturbations as well as robustness to actuator failure will provide important improvements to the MAV control problem.

Acknowledgements

This work was partially supported by the Air Force Office of Scientific Research (AFOSR) grant number FA9550-08-1-0187 and the National Science Foundation (NSF) grant IIS-0910358.

References

- [1] M. Abdulrahim and J. Cocquyt. Development of mission capable flexible-wing micro air vehicle. In *53rd Southeastern Regional Student Conference*, 2002.
- [2] M. Abdulrahim and R. Lind. Investigating segmented trailing-edge surfaces for full authority control of a uav. In *AIAA Atmospheric Flight Mechanics Conference*, 2003.
- [3] A. Agogino and K. Tumer. Efficient evaluation functions for multi-rover systems. In *The Genetic and Evolutionary Computation Conference*, pages 1–12, Seattle, WA, June 2004.
- [4] A. Agogino and K. Tumer. QUICR-learning for multi-agent coordination. In *Proceedings of the 21st National Conference on Artificial Intelligence*, Boston, MA, July 2006.
- [5] Richard K. Arning and Stefan Sassen. Flight control of micro aerial vehicles. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2004.
- [6] S. R. Bieniawski. *Distributed Optimization and Flight Control Using Collectives*. PhD thesis, Stanford University, 2005.

- [7] S. R. Bieniawski, I. Kroo, and D. Wolpert. Flight control with distributed effectors. In *AIAA Guidance, Navigation, and Control Conference, San Francisco, CA, August 15-18, 2005*.
- [8] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, November 1995.
- [9] S. Farritor and S. Dubowsky. Planning methodology for planetary robotic exploration. In *ASME Journal of Dynamic Systems, Measurement and Control*, volume 124, pages 4: 698–701, 2002.
- [10] H. Garcia, M. Abdulrahim, and R. Lind. Roll control for a micro air vehicle using active wing morphing. In *AIAA Guidance, Navigation and Control Conference*, 2003.
- [11] F. Gomez and R. Miikkulainen. Active guidance for a finless rocket through neuroevolution. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Chicago, Illinois, 2003.
- [12] J. Hall, D. Lawrence, and K. Mohseni. Lateral control and observation of a micro aerial vehicle. In *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007.
- [13] David A. Jenkins, Peter G. Ifju, Mujahid Abdulrahim, and Scott Olipra. Assessment of controllability of micro air vehicles. Technical report, University of Florida, 2000.
- [14] T. Kordes, M. Buschmann, S. Winkler, H.-W. Schulz, and P. Vorsmann. Progresses in the development of the fully autonomous MAV CAROLO. In *2nd AIAA Unmanned Unlimited Systems, Technologies, and Operations Aerospace*, 2003.
- [15] Roman Krashanitsa, George Platanitis, Bill Silin, and Sergey Shkarayev. Aerodynamics and controls design for autonomous micro air vehicles. In *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, 2006.
- [16] M. A. Marra, B. E. Boling, and B. L. Walcott. Genetic control of a ball beam system. In *IEEE international Conference on Control Applications*, 1996.
- [17] David Moriarty and Risto Miikkulainen. Forming neural networks through efficient and adaptive coevolution. *Evolutionary Computation*, 5:373–399, 2002.
- [18] Nikhil Nigam and Ilan Kroo. Control and design of multiple unmanned air vehicles for a persistent surveillance task. In *AIAA*, 2008.
- [19] Nikhil Nigam and Ilan Kroo. Persistent surveillance using multiple unmanned air vehicles. *IEEE*, 2008.
- [20] L. Panait, S. Luke, and R. P. Wiegand. Biasing coevolutionary search for optimal multi-agent behaviors. *IEEE Transactions on Evolutionary Computation*, 10:629–645, 2006.
- [21] F. Pasemann. Evolving neurocontrollers for balancing an inverted pendulum, 1998.
- [22] William J Pisano, Dale A Lawrence, and Peter C Gray. Autonomous uav control using a 3-sensor autopilot. In *AIAA Conference and Exhibit*, 2007.
- [23] M. A. Potter and K. A. de Jong. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1):1–29, 2000.

- [24] P. Scerri, R. Glington, S. Owens, D. Scerri, and K. Sycara. Geolocation of RF emitters by many uavs. In *AIAA Infotech@Aerospace 2007 Conference and Exhibit*, 2007.
- [25] Sandip Sen, Sandip Debnath, and Manisha Mundhe. Evolving coordinated agents. *Advances in evolutionary computing: theory and applications*, pages 559–577, 2003.
- [26] K. Stewart, J. Wagener, G. Abate, and Max Salichon. Design of the air force research laboratory micro aerial vehicle research configuration. In *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007.
- [27] P. Jan 't Hoen and E. de Jong. Evolutionary multi-agent systems. In *Proceedings of the 8th International Conference on Parallel Problem Solving from Nature*, pages 872–881, 2004.
- [28] K. Tumer and A. Agogino. Coordinating multi-rover systems: Evaluation functions for dynamic and noisy environments. In *The Genetic and Evolutionary Computation Conference*, Washington, DC, June 2005.
- [29] K. Tumer and A. Agogino. Distributed agent-based air traffic flow management. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 330–337, Honolulu, HI, May 2007.
- [30] K. Tumer and D. Wolpert, editors. *Collectives and the Design of Complex Systems*. Springer, New York, 2004.
- [31] Kagan Tumer and Adrian Agogino. Distributed evaluation functions for fault tolerant multirover systems. In *Genetic and Evolutionary Computation Conference*, 2006.
- [32] Martin R. Waszak, John B. Davidson, and Peter G. Ifju. Simulation and flight control of an aeroelastic fixed wing micro aerial vehicle. In *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, 2002.
- [33] Martin R. Waszak, Luther N. Jenkins, and Peter Ifju. Stability and control properties of an aeroelastic fixed wing micro aerial vehicle. In *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, 2001.
- [34] D. H. Wolpert and K. Tumer. Optimal payoff functions for members of collectives. *Advances in Complex Systems*, 4(2/3):265–279, 2001.